

Claims

What is claimed is:

1. A processor comprising:

scheduling circuitry operative to schedule data blocks for transmission from a plurality of transmission elements, the scheduling circuitry being configurable for utilization of at least a first table and a second table in scheduling the data blocks for transmission; and

memory circuitry associated with the scheduling circuitry and configurable to store at least a portion of at least one of the first and second tables;

the first table configurable to include at least first and second lists of entries corresponding to transmission elements for which data blocks are to be scheduled in accordance with at least a first scheduling algorithm, the scheduler being operative to maintain a first table pointer identifying at least one of the first and second lists of the first table as having priority over the other of the first and second lists of the first table;

the second table configurable to include a plurality of entries corresponding to transmission elements for which data blocks are to be scheduled in accordance with at least a second scheduling algorithm different than the first scheduling algorithm, association of a given one of the transmission elements with a particular one of the entries establishing a scheduling rate for that transmission element, the scheduler maintaining a second table pointer identifying a current one of the second table entries as being eligible for transmission.

2. The processor of claim 1 wherein the first table comprises a plurality of first-in first-out lists, the plurality of lists including the first and second lists and at least one additional list.

3. The processor of claim 1 wherein the first table pointer comprises an active list pointer identifying one of the first and second lists as an active list.

4. The processor of claim 1 wherein in scheduling data blocks for transmission utilizing the first table, the scheduling circuitry identifies a first entry in a first non-empty one of the lists starting

from a list identified by the first table pointer, and schedules for transmission a data block from the corresponding transmission element.

5 5. The processor of claim 1 wherein one of the first and second lists of the first table
comprises an active list and the other of the first and second lists of the first table comprises a
pending list, wherein the active list comprises a list of one or more of the transmission elements
which have not exceeded corresponding bandwidth allocations, and the pending list comprises a list
of one or more of the transmission elements which have exceeded corresponding bandwidth
allocations, the bandwidth allocations being determined over a programmable time interval based
10 on relative rate partitions between the transmission elements.

15 6. The processor of claim 5 wherein the first table pointer initially points to the active list,
and when the active list becomes empty, the pointer is updated to point to the pending list, the
pending list is designated as a new active list, and the previous active list is designated as a new
pending list.

20 7. The processor of claim 5 wherein a given transmission element which has exceeded its
corresponding bandwidth allocation is moved from the first table to the second table in order to
provide an adjustment in its scheduling rate.

25 8. The processor of claim 1 wherein the second table includes a plurality of slots, with each
slot capable of storing the identifier of one of the transmission elements.

9. The processor of claim 1 wherein the second table pointer comprises a current pointer
which is incremented in accordance with a time base of the scheduling circuitry.

10. The processor of claim 1 wherein the second table is configured such that each of the
plurality of transmission elements can be dynamically assigned a transmission rate.

11. The processor of claim 1 wherein a desired scheduling rate for a given one of the transmission elements is established by entering an identifier of that element into a particular one of the entries of the second table, the particular entry being determined as a function of the second table pointer and a designated scheduling interval.

5

12. The processor of claim 11 wherein the designated scheduling interval is dynamically alterable under software control.

10 13. The processor of claim 1 wherein each of the transmission elements at a given point in time may have a corresponding entry in the first table or in the second table, but not in both the first table and the second table.

15 14. The processor of claim 1 wherein the first scheduling algorithm comprises a weighted fair queuing scheduling algorithm, the weighted fair queuing algorithm being configurable to implement at least one of round robin scheduling, strict priority scheduling, weighted round robin scheduling, smooth weighted round robin scheduling and smooth deficit weighted round robin scheduling.

20 15. The processor of claim 1 wherein the second scheduling algorithm comprises at least one of a constant bit rate scheduling algorithm and a variable bit rate scheduling algorithm.

16. The processor of claim 1 wherein the memory circuitry comprises at least one of internal memory and external memory of the processor.

25 17. The processor of claim 1 wherein one or more of the data blocks comprise data packets.

18. The processor of claim 1 wherein the processor comprises a network processor integrated circuit configured to provide an interface for data block transfer between a network and a switch fabric.

19. A method for use in a processor, the method comprising:

- storing at least a portion of at least one of a first table and a second table; and
- scheduling data blocks for transmission from a plurality of transmission elements, utilizing the first and second tables;
- the first table configurable to include at least first and second lists of entries corresponding to transmission elements for which data blocks are to be scheduled in accordance with at least a first scheduling algorithm;
- a first table pointer identifying at least one of the first and second lists of the first table as having priority over the other of the first and second lists of the first table;
- the second table configurable to include a plurality of entries corresponding to transmission elements for which data blocks are to be scheduled in accordance with at least a second scheduling algorithm different than the first scheduling algorithm, association of a given one of the transmission elements with a particular one of the entries establishing a scheduling rate for that transmission element;
- a second table pointer identifying a current one of the second table entries as being eligible for transmission.

20. An article of manufacture comprising a machine-readable storage medium for use in conjunction with a processor, the medium storing one or more software programs for use in scheduling data blocks for transmission from a plurality of transmission elements, the one or more programs when executed implementing the step of:

scheduling data blocks for transmission from a plurality of transmission elements, utilizing the first and second tables;

the first table configurable to include at least first and second lists of entries corresponding to transmission elements for which data blocks are to be scheduled in accordance with at least a first scheduling algorithm;

5 a first table pointer identifying at least one of the first and second lists of the first table as having priority over the other of the first and second lists of the first table;

the second table configurable to include a plurality of entries corresponding to transmission elements for which data blocks are to be scheduled in accordance with at least a second scheduling algorithm different than the first scheduling algorithm, association of a given one of the transmission elements with a particular one of the entries establishing a scheduling rate for that
10 transmission element;

a second table pointer identifying a current one of the second table entries as being eligible for transmission.